

SAM (REX/Ops) Operations Meeting 2007.11.27

Present: Adam (typist), SteveW, Andrew, Art, Robert, Dave, Dehong, Krzysztof, Parag, Rick St. D, SteveS, Angela, Anil, Randolph

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

Randolph will try to do the Retired Date testing today since there's a short gap in the NGOP stuff.

Andrew has a new station that fixes the project manager memory leak (see D0 section below).

MINOS needs their sam_products.

DØ:

Events (worth noting) :

-- Shifters reported 4 (enstore) movers offline since ~Sat.

-- d0cabsrv2 was hung ~ last Thursday fnal-cabsrv2 was dead at that time.

On Tuesday d0cabsrv1 crashed with out of memory. On Thursday morning the same thing happened to d0cabsrv2. Have noticed very high memory usage on d0cabsrv1. Problem was many project masters with large memory usage. And we have > 250 projects running. Robert ran through valgrind - found memory leaks in base code (CDF has these too) and MIS events. Andrew has released a new project master that fixes the memory leaks. Robert put this on Clued0 on Wednesday. Nothing bad happened over the long weekend. It is now installed on CAB. Now it looks like the PM's are staying below 15 MB.

There is a small memory leak every time the station gets a new proxy from the DB server. But this is only a few 100 bytes, so would only affect extremely long running stations. And could cause memory fragmentation. Should be fixed in next station version.

There are 4 GB of memory on the machines. Will be transferring to new machines soon.

-- Some issues with fnal-cabsrv1 (eg. down a few hours last Wed.) earlier this week and a lot other issues with sam-grid as reflected in the tickets below.

See above.

Quite a few unresolved issues (probably related to the holidays):

#3177/3178: "Could not get sam incoming storage name" which might be related to an AFS problem and more ??

Remote site problems (probably).

#3176: "job submission problem at grid1.oscer.ou.edu". Please help.

Parag will look into this.

#3173: "Files cannot be declared at NIKHEF" -> MC jobs stalled after 250 events. Please help.

UK dCache problem?

#3171: "samgrid job error: Globus error 155" for "data reconstruction" jobs. Please help.

#3169: "Problems with merge jobs in stuck in idle state" : waiting to be tested out.

#3158 : "MSU site not appearing in list of resources" : seen some long exchanges between Tom and Parag but still resolved

#3147 : "Unable to get dataset at UNL" --- updates from Parag and/or Andrii ?

#3146: "jobs to manchester being held" -- famous "Globus error 24" problem ?

CDF:

- o SAMTV fix did not work -- cycled once and failed

Angela is trying to figure this out.

- o r-fcdfdata320 inactive (still)

- o Upgrade of LT03 firmware finished at 11:29 last Tuesday. Tapes worked fine through the week.

- o One of the stken servers, stkensrv3, died shortly before midnight Tuesday. Although it has important duties in enstore, it is not vital, and the public systems continued to function while it was down. Dave Berg fixed it by 1:30am on Weds.

- o ENCP tape error noted by shift crew by 4:20 on 11/21. Dehong noted by 10am that it was ok, there was no error. We had discussed this at the SAM meeting last Tuesday.

Krzysztof asked Ray Culbertson to add some delay to the reporting script. He also pointed out that the graph of transfer rates had lots of empty spaces so there was plenty of capacity.

- o The movers:

- LTO3_01

- LTO3_26

- LTO3_29

- died and were restarted on Wed.

- o Krzysztof noted some dcache mover queuing, so decided to move several pools to TestPools and to increase the max mover parameters to drain the backlog. He kept an eye on it. Pools moved to TestPools are listed below.

- ReadOnlyPools

- r-fcdfdata145-1

- r-fcdfdata145-2

- r-fcdfdata145-3

- r-fcdfdata154-3

- RawPools

- r-fcdfdata124-1 mover set max active 24

- r-fcdfdata124-2 mover set max active 24

- r-fcdfdata124-3 mover set max active 24

- r-fcdfdata327-1 mover set max active 60

Without Krzysztof's intervention, dCache probably would not have recovered (pools would have crashed).

- o Thurs night: one bad mover

- LTO3_04. Sent mail and it was fine on Friday.

- o 1am on Sat morning: 3 bad drives and a bad tape

- LTO3_01

- LTO3_28

- LTO3_30

- IAB444 0.45GB (NOACCESS 1123-2201 full 0522-1800) CDF-LTO3

cdf.cdfpewk.cpio_odc

- By morning the bad tapes drive were no longer bad but the tape was still bad. Sent mail Sat morning. It was fixed Monday.

- o Monday Morning, two 9940B tapes went no-access. Sent mail on this and repeated request on IAB444.

- IA8992 0.18GB (NOACCESS 1125-1922 full 0909-0349) CDF-9940B

cdf.xpmm0h.cpio_odc

- IAA641 0.86GB (NOACCESS 1125-1918 full 0804-1152) CDF-9940B

cdf.cdfpstn.cpio_odc

- All tapes were ok by 1pm.

MINOS:

Minos status 2007 Nov 27

General -

AFS - Servers upgraded for stability Saturday 06:00 to 11:00
Timeouts of some servers continued Monday.

Enstore -

DCache -

BlueArc -

Need AFS/NFS translated or remounted mount of products and releases on Farm.

SAM -

IT 3128. sam_products for sam_station v6_0_5_22_srm
v4_29 describes sam_station v6_0_3_9p1 -q GCC-3.1

IT 1894 Need 'sam locate' in C++ API

Oracle -

OPEN ISSUES

IT 1146: cannot test existence of sam tape location
IT 1894: sam locate via C++ api, for simple client needs
IT 1979: Compiler warnings from sam_cpp_api
IT 2071: quadratic project startup time growth
IT 3128. sam_products for sam_station v6_0_5_22_srm

PLAN -

Migration to 9940B tape can proceed anytime

DB:

D0 :

- D0ora1 will be down starting at 8am, Tuesday 11/27, for OS patching. Downtime is expected to be 4 hours.

- Dec 04, d0ora2 OS patches and remove the old clariion array.
Long live the Clarion array!

- SAM cut v6_10 on d0ofdev1 - To add composite index on lower(file_name) and retired_date -done

- v6_4 index on create_date of data_files against d0ofdev1 done
- Service Request with Oracle on Full table scan Vs Index scan on data_files_param_values.
Development is working on Bug Fix. Bug# 5601578 No time estimate yet for bug fix.
SAR# 5837876.993

Alerts:

CDF :

- DB patches on cdfdev/cdfint Nov 27 8:00AM-11:00AM
- DB patches fcdfora4~cdfopr2 1HR } Dec 04
O/S Qrtly fcdfora4~cdfopr2 2HRS }

Alerts:

MINOS:

Anil and Art will follow up about downtime to install patches.

enstore/dCache:

enstore/dcache report

Scheduled Downtimes:

- o D0 - Tue, Dec 4
No plans (yet)
- o CDF dCache - Tue, Dec 4
 - If CDF decides to take a downtime, dCache will be restarted.

Operations:

- o SL8500s at GCC
 - Firmware upgrade completed by about 11:30 Tuesday.
 - Both SL8500s and the pass-thru are now active.
 - LTO4 drives remain to be moved from FCC, will not require downtime.

- Some tapes with our own "wide" labels were not reliably recognized by the new firmware. To date, 100 CMS tapes and 1 CDF tape needed their labels replaced.

o d0en

- Some LTO2 blanks were made writeable today, more to follow.

o cdfen

- One LTO3 mover died but rebooted after power cycled, another apparently lost network for a few hours, but recovered on its own.
- The raid disk reported last week had failed on Sun, Nov 11.

o cdfdca (from Timur's report)

- Krzysztof found that some of the pools are queuing movers and moved the pools to TestPools pool group to allow the draining and reconfiguration.
- Mover queuing is a normal dCache feature and unless the pool selection is highly unbalanced, it should not be considered a problem.
- Timur notified cdf about the env variable to control the active/passive mode of the dcap client transfer. Krzysztof requested that we make server passive mode a default in kits. We need to discuss the impact of this change first before we agree to implement it. I am concerned about other users who use dccp from kits.

o stken

- stkensrv3 (backups, enstore/dcache monitoring) was replaced. There was no effect on enstore operations while it was down.

We appreciate Dave coming in at ~1am to replace this machine.

o fndca

- Pools on fndca9a remain out of the configuration. One raid controller has been replaced and is undergoing tests.